Method of synchronizing message transmissions

The invention relates to a method of synchronizing message transmissions between mobile nodes in an ad-hoc network using a medium access protocol.

In an ad-hoc network of mobile nodes which organizes itself for the purpose of exchanging information between the nodes without a stationary control device, the messages are exchanged between the individual nodes in a wireless manner via radio links. Such an ad-hoc network is formed, inter alia, by motor vehicles in road traffic which are equipped with the appropriate transmitters/receivers and other devices known to the person skilled in the art for evaluating and creating desired messages. Each of the motor vehicles may in this case be regarded as a mobile node of the ad-hoc network. Using such networks, it is possible for example for an emergency vehicle, when approaching a crossroads, to make the other motor vehicles aware of its presence by exchanging corresponding information and for the crossroads to be kept free in order for the emergency vehicle to pass quickly through. The motor vehicles may likewise exchange information about their respective speeds, in order to avoid a collision when changing lanes or merging into a lane.

For this purpose, use is made inter alia of medium access protocols such as the IEE 802.11 protocol, which regulates access to a common communication medium without the use of a central control device. The RTS and CTS parts of the MAC protocol are very useful in order to overcome the "hidden node problem" and to avoid overlapping in message transmission. The hidden node problem arises when two nodes approach an intermediate third node from different directions, wherein the transmitting/receiving ranges of the two approaching nodes do not yet overlap so that these two nodes cannot perceive one another. The RTS and CTS parts of the protocol are suitable for a point-to-point communication but not for a point-to-multipoint communication. Point-to-point communication is also referred to as "unicast" operation and point-to-multipoint communication is also referred to as "broadcast/multicast" operation. A similar problem arises in respect of confirming correct receipt of a message. A peer-to-peer connection is to be understood as meaning a communication between one transmitter and one receiver. The transmitter firstly transmits an RTS (request to send) signal which comprises both the address of the transmitter and that of the receiver. The receiver then transmits a CTS (Clear To Send) signal comprising the

address of the transmitter. As a result, transmitter and receiver are in each case prepared for a subsequent data transmission between them. The other nodes of the network overhear the RTS and CTS signals and will accordingly themselves not pick up any communication for the duration transmitted with these signals. In this case, the hidden node problem is solved in that the CTS signals sent out by the receiver are also received by nodes which lie out with the range of the transmitter. If the receiver were not to transmit a CTS signal, the transmitter would also not transmit any data and would start a new attempt by sending out a new RTS signal.

Message transmission in "broadcast/multicast" operation is susceptible to instances of overlapping of message transmissions and to incorrect receipt of messages, and therefore has only low capability. In a wireless ad-hoc network without a central control device, the number of participating nodes and the number of connections between the nodes change continuously. Moreover, the external circumstances change continuously.

From the publication Gökhan Mergen and Lars Tong: "Receiver controlled medium access in multihop ad hoc networks with multipacket reception", Military Communications Conference, 2001 MILCOM 2001, Pages 1014-1018 vol. 2, 28-31 Oct. 2001 (see part 1-3,2) [20011028 Mergen] and WO 02/15485 A 1, a method is known in which the receivers control message transmission. In this case, however, there is only a single receiver per transmitter, so that multicast operation is not supported.

In the publication Lars Wischhof et al: "Adaptive broadcast for travel and traffic information distribution based on inter-vehicle communication [20030609 Wischhof Lars], a method for self-organization of an ad-hoc network is described, wherein the "transmission interval" of each node can be adapted by each node to the local environment in order to make efficient use of the available bandwidth. However, an adaptation of the "transmission interval" will disrupt a potential synchronization of the nodes with one another, as a result of which potential overlaps can less easily be foreseen and therefore can rarely be avoided.

In the publication André Ebner et al: "Decentralized slot synchronization in highly dynamic ad hoc networks", Proceedings of the 5th International Symposium on Wireless Personal Multimedia Communication (WPMC'02) Honolulu Hawaii, Oct. 27-30, 2002 (see pages 1-3) [20021027 Ebner Andre], a decentralized time slot synchronization is proposed, wherein the local time slot of a node is shifted in accordance with the received

time slots by a certain fraction of these time slots. In this case, there is an idle time between the individual time slots in order to take account of the different lengths of the time slots. It is nevertheless possible for this idle time to be insufficient to compensate the difference between the (local and received) time slots and this accordingly leads to overlaps or message losses.

US 6,565,582 B1 discloses a method for message transmission, wherein a transmitter and a receiver communicate with one another in the handshake mode and thus a third transmitter/receiver is excluded from access to this transmission channel. This method, however, does not support multicast operation.

WO 02/28020 A2 discloses a method for avoiding overlaps during message transmission via a node, wherein it is determined by means of known algorithms which node is sending in which time slot.

It is an object of the invention to specify a method for message transmissions, by means of which a synchronization of the individual nodes with one another is possible without the use of a central control device. In particular, a time division multiplexing method is to be provided.

This object is achieved by the features specified in claim 1.

The core concept of the invention is that the messages which are sent and received by the individual nodes of the ad-hoc network are compared with one another in order to make it possible to predict the time of sending of the messages. Furthermore, the nodes are synchronized with one another such that they do not collide with one another or overlap when one of the nodes perceives that other nodes are in its transmitting/receiving range. In addition, a node which receives two messages that collide with one another or two messages that overlap from two other nodes reports this to these two nodes. It is understood that, in the case of three or more overlapping messages, all sending nodes are accordingly informed by the receiving node.

Within the context of the invention, the comparing of the messages with one another and the synchronization of the nodes can be carried out by the person skilled in the art in any desired manner, but are preferably carried out as described below.

The advantage of the invention is that there is no need for a central control device which synchronizes the exchange of messages between nodes of the ad-hoc network by means of an internal clock in the control device, so that in particular a network of mobile

nodes, as occurs inter alia in road traffic, can organize itself in order to obtain a network with a maximum throughput of information which is exchanged between the individual nodes or motor vehicles of the network. A stable and reliably operating network can thus likewise be provided.

Advantageous refinements of the invention are characterized in the dependent claims.

Preferably, in accordance with claim 2, a TDMA-type protocol known to the person skilled in the art is used to exchange messages between the nodes in a time division multiplexing method without the use of a central synchronization clock. In a TDMA protocol, the time slots in a data transmission block or "frame" determine the connections and the times at which the transmitters are to transmit their data. In order to allow this fixed access, the messages that are to be transmitted are compared with one another in the following manner: Each message must have the same length. If a longer message is to be transmitted, this must be split into a number of segments; a shorter message is brought to the appropriate length by means of "padding bits" in the manner of a buffer signal or by means of empty filling bits. Furthermore, each transmitter transmits its messages with a message transmission rate that is constant for the transmitter. If no messages are to be transmitted by the transmitter, empty messages may be transmitted. This can be avoided in that in this case, e.g. the last status of the vehicle is transmitted, that is to say inter alia its location and its speed. In addition, each transmitter/receiver checks the transmission medium to ascertain whether it is free at that moment, that is to say no other data or messages are being transmitted. Access to the medium is controlled in accordance with the DCF part of the 802.11 MAC protocol. Following successful transmission of a message, the node waits for a certain period of time, which depends on the message transmission rate, before it again attempts to access the medium in accordance with the DCF protocol.

If all nodes were to transmit with the same message transmission rate, a synchronization of all nodes would automatically be achieved since each node would then send only when no other nodes are sending. In the case of a mobile ad-hoc network, however, this is much more complex since a newly added node is not synchronized with the existing communication and does not necessarily transmit its messages with the same message transmission rate and thus can disrupt the existing synchronization.

With a limitation of the message transmission rate for each node, as specified in claim 3, a synchronization of the nodes of the network can be achieved in particular when the number of nodes is small enough to have in the available bandwidth of the radio

frequencies of the ad-hoc network enough free capacity that the message transmission rates can fluctuate.

Although the need for each node to send with a message transmission rate determined for it is not a flexible solution, it may be advantageous for some situations. As specified in claim 4, a node sends a presence signal in order for example to make the other road users aware of the presence of this motor vehicle. By way of example, such a presence signal is sent on a specific radio frequency on which no other messages must be exchanged. The presence signal may inter alia contain an identification code, a position information item, a speed and other information such as the type of motor vehicle. The traffic density can also be deduced from these presence signals. This may be calculated for example from all the presence signals and the positions. In the event of an accident, corresponding warning signals would be sent out by a large number of motor vehicles in order to make the other road users aware of the accident. In this case, the motor vehicles or nodes should in each case have available only a limited message transmission rate in order to avoid overlaps between too many messages sent out and a collapse of the network. In order to solve this problem, the presence signal and a warning signal may be combined with one another in order to form a message of specific length which requires only a maximum transmission rate. In a normal traffic situation, this message of specific length may consist almost entirely of information about the motor vehicle or other information that is to be exchanged, without the network being overloaded. In the event of an accident, however, the presence signal can be transmitted with maximum transmission rate in order to warn other road users.

In order to optimize message transmission between the individual nodes, each has an individual transmission rate as described in claim 5. This may be achieved in that a frame of message slots is defined. The frame size, that is to say the number of message slots, is defined for the entire network. Each node defines the start of its own frame. The message transmission rate of this node is then defined by the number of message slots occupied by this node. If sending is to be carried out with only a low rate, for example, only one of the message slots is occupied by this node. In the case of a high transmission rate, a number of message slots may be occupied.

In order to synchronize message transmission between a number of nodes, a frame of a node may be temporally shifted in order to send at times at which the other nodes in the ad-hoc network are not sending, as specified in claim 6. If a motor vehicle is approaching an existing ad-hoc network, this may lead to overlaps in message transmission, since this new node is not yet synchronized. Synchronization can be brought about by using

the IEEE 802. DCF protocol. Such a synchronization is particularly suitable when the node uses only a single slot of a frame and does not occupy any slot of another node. The new node finds a slot in a simple manner by waiting long enough until the medium is free, that is to say none of the other nodes are sending a message. If a number of slots are occupied by the new node, those slots which are already occupied by other nodes have to be made free and the message transmission rate has to be reduced at least temporarily. A node which ascertains that a slot is occupied by another node assumes that it itself is not synchronized and terminates the transmission of messages in this slot and sends only in the remaining slots occupied by it and not by other nodes, with a lower message transmission rate. In order to pass back again to the full message transmission rate, a complete frame is monitored in order to discover the free slots and use them for its own message transmission.

Given a typical bandwidth of an available radio frequency of 20 Mb/s and for example 400 nodes or motor vehicles in the ad-hoc network, these may transmit a radio message with 1000 bits every 50 ms. This corresponds to a 40% use of the bandwidth. The duration of a frame would in this case be 50 ms. If in this case two nodes approach one another with a relative speed of 100 m/s and a perception range of 100 meters, a distance of 90 to 95 meters is available to them once they have synchronized with one another, if no overlapping of newly occupied slots occurs. Accordingly, they may communicate with one another for approximately 38 frames before they move out of range. If a large number of nodes are moving in opposite directions, such as motor vehicles on a country road, practically no synchronization can be produced between them. In this case it is recommended that the nodes or motor vehicles moving in different directions communicate with one another in each case on different radio channels, so that only the nodes moving in one direction form a network.

In the IEEE 8002.11 protocol, the hidden node problem is solved by means of RTS and CTS signals. A receiving node will not transmit any CTS signal if two or more RTS signals directed at it overlap one another, so that the sending nodes do not transmit any data but rather try again, in a manner offset in time, to establish a communication. In multicast operation, the RTS and CTS protocol is of little use, since a handshake mode has to be carried out with each receiver one after the other. Instead of transmitting a CTS signal to the nodes from which an RTS signal has been received without error, it is proposed that the receiving node carries out quite the opposite, namely informs the transmitting nodes that they should not send. Since the messages are transmitted in message slots in frames, the receiving node which ascertains an overlap can inform the sending nodes about the overlap in the next

frame. This may be effected in any desired manner. The basic idea is that in one frame an overlap is ascertained in a message slot and this is informed to the nodes in the next frame. Accordingly, in the next frame they can in turn select a free slot for transmitting messages.

Normally, receipt of a message by the receiving node is confirmed by a confirmation report only in the case of a direct connection or message transmission between two nodes. If messages are transmitted by one node to a large number of other nodes, such a confirmation does not take place, since it would take too long to receive a large number of confirmation reports, so that no effective communication would be possible. With a TDMA-type organization of the message transmission in frames with a defined number of n (time) slots, a possibility is given to transmit such confirmations.

This consists in that a vector is composed of confirmation elements, wherein each confirmation element is assigned to a slot of the frame.

Each node which transmits a message in a slot also transmits a confirmation vector with (n-1) bits, wherein the vector represents confirmation of the correct receipt of a message in one of the preceding (n-1) slots. It is thus possible for each node to check whether all the other nodes involved in the communication have correctly received its message within a frame and each node can check which slots are occupied. This cannot only be effected by direct monitoring of the transmissions of messages in this slot but rather also indirectly by interpreting the received confirmation vector. If any of the received vectors indicates that a message has not been correctly received, for example in a slot k, this slot is marked as occupied by the receiving node even if the receiving node itself has not received any message from this other node in this slot. In this case, each node monitors whether and, if yes, which slots are occupied by messages and which are free.

A node which wishes to transmit a message in a new slot monitors the exchange of messages for a period of the frame and the associated confirmation vectors which are stored for example in a suitable memory. From this information, it can directly and/or indirectly select a free slot, provided one is available, in which no messages have yet been transmitted. In this slot, in the next frame it transmits its first message and also a confirmation vector during the next period of data transmission.

A node which has transmitted a message in the preceding frame implicitly reserves the same slot for the next frame, as in an R-ALOHA system known to the person skilled in the art.

Before the node transmits a message in a slot, it checks whether it has received the corresponding confirmation reports from all the other nodes which have transmitted a

message within the time between this slot in the preceding frame and the slot in the new frame. If one or more confirmation reports are missing, this is stored by the node. If in the next frame one or more of the already missing confirmation reports are once again not received, this is compared with the stored missing confirmation reports and from this it is determined whether a corresponding confirmation report is missing from one or more nodes in the two preceding periods of data transmission. Accordingly, the receiving node marks this slot as occupied by another node, and the node does not use this slot for its own message transmission but rather selects another slot.

A similar procedure is described in the publication F. Borgonovo, A. Capone, M. Cesana, L. Fratta: "ADHOC MAC: a new, flexible and reliable MAC architecture for ad-hoc networks", University of Mailand. In this case, an information item is transmitted by each node in order to indicate the occupied and free slots in a frame. A slot is marked as occupied and shown in the following information item if the node has received a message correctly and in full. This corresponds to the confirmation vector described above.

If a node receives a message from two or more other nodes in the same slot of the frame, the received messages usually differ in their reception strengths on account of the different distances of the nodes from one another, so that only the message with the greatest reception strength is classified as correctly received by the node. In this case, the receiving node sends a confirmation report for this slot. In this publication, an identification signal relating to the node from which the message has been correctly received is also transmitted by the receiving node. The nodes sending in the same slot can thus in each case clearly ascertain whether the receiving node has output a confirmation report for the message transmitted by it, although this is only the case for the one sending node with the correctly received message.

By including the identification of the transmitting node in the confirmation report, however, the amount of information to be transmitted is undesirably increased, so that the capacity for message transmission is reduced. This is the case in particular when a large number of slots are provided in a frame and/or when each node has an identification code that is unique throughout the world, which requires a large number of bits for encryption.

Accordingly, this identification code is omitted in the confirmation vector, so that this vector comprises precisely (n-1) bits. Admittedly, in principle, there is the possibility in this case that some nodes which have transmitted messages will receive a confirmation signal that was intended for another node. Although this would make the method of

confirming correct receipt of a message less reliable, the hidden node problem is nevertheless thus reliably overcome.

By contrast, in the abovementioned publication, in the case of an incorrect message transmission, for example in the case where not all the received confirmation vectors indicate occupation of a specific slot, this slot is released again in the next frame for message transmission. According to the invention, this is avoided in that a slot is released only once no confirmation vectors have been received by a node in two successive frames for a specific slot. By omitting an identification code for the correctly sending node, the number of data to be transmitted by the receiving node is reduced, in particular the confirmation vector has only (n-1) bits in the case of n slots in the frame, instead of at least n elements.

It will be understood that the method of synchronizing message transmissions in an ad-hoc network for self-organization of the nodes can be used in any systems. By way of example, the containers in a container terminal may be equipped with the appropriate devices to form a node, so that the goods flow of the containers can be optimized. Preferably, as specified in claim 8, the method is used to control a flow of traffic, in order that the transmission of messages between the individual road users or motor vehicles is optimized and in particular warnings are reliably transmitted to the other road users. In the ad-hoc network, any desired data can be exchanged, e.g. telephone conversations between the road users, as long as a connection exists between the nodes.

The invention will be further described with reference to examples of embodiments shown in the drawings to which, however, the invention is not restricted.

Fig. 1 shows an example of the hidden node problem.

Fig. 2 shows a frame with message slots.

Fig. 3 shows the synchronization of message transmission.

Fig. 4 shows overlapping messages.

Fig. 5 shows overlapping messages.

Fig. 1 shows three nodes A, B, C with their respective radio ranges r, wherein nodes A and C are moving toward one another in the direction of the resting node B, as shown by the arrows. The distance between the two nodes A and C is so great that node A cannot receive any signals from node C and vice versa. It is therefore possible for the two to

transmit a message at the same time which is received by node B. Such an overlap is referred to as the "hidden node problem". If such an overlap is ascertained by node B, a corresponding signal is transmitted to nodes A and C, which then use another slot for message transmission.

Fig. 2 shows successive frames which in each case consist of n message slots. A node in this case uses two message slots, namely the first and the (n-1)th, to transmit messages to other nodes, as shown by the vertical arrows. The horizontal arrow shows the length of a frame.

An automatic synchronization is shown in the upper half of Fig. 3, this taking place when a new node is added to an existing network. Node 2, which is already present in the network, sends in its frame in slots 1 and n-1, the start of the frames of nodes 1 and 2 being at different times. The newly added node 1 monitors the medium until no message is being sent. The frame of node 1 is thus shifted until its slot 4, in which it would like to transmit a message, is arranged in time terms such that it lies after slot 1 of node 2, as shown on the right-hand side.

An alternative procedure of the new node 1 is shown in the lower half of Fig. 3. Originally, the newly added node 1 wishes to transmit a message in its frame in slots 2 and 4. When first monitoring the message transmissions of the other nodes, it ascertains that at the time of its slot 4 the medium is occupied by node 2. Node 1 then shifts its message transmission from its slot 4 to its free slot 5, which corresponds to slot 3 of the frame of node 2, as shown on the right-hand side.

Fig. 4 shows an example for handling the hidden node problem. Node B ascertains an overlap of message transmissions in its slot 0, wherein node A has transmitted a message in its slot 2 and node C has transmitted a message in its slot 6. Node B transmits a message in its slot 3. Once it has ascertained this overlap, node B transmits a corresponding information item to nodes A and C, which may also contain in each case a free slot for these nodes. The manner of transmitting this notification can be configured as desired. Thereafter, as can be seen in the right-hand half of Fig. 4, node A transmits a message in its slot 4 instead of slot 2 and node C transmits a message in its slot 7 instead of slot 6, with no overlap.

An alternative possibility consists in that, as shown in Fig. 5, node B decides no longer to transmit its messages in its slot 3 but rather to transmit them in its slot 0. This in turn is ascertained by nodes A and C, whereupon these then select in each case a new slot for message transmission. Changing of the message transmission slot by node B is understood by nodes A and C as a signal that their messages have overlapped.

In order to reduce the likelihood that in the event of a new slot occupation two or more nodes will use the same message slot for message transmission, the slots may be randomly selected from the freely available slots. Furthermore, it is possible for each node to register the occupied message slots in temporal sequence. For this, each node carries an occupation list, in which for each slot an occupation in the preceding frame is stored. If a slot is not occupied, the fact of which node occupied this slot however many frames ago is likewise registered. A corresponding example is shown in the table below:

slot number	occupied	node number	how often not occupied
0	yes	123	
1	no	222	23
2	yes	342	
3	no	143	20
4	no	342	1

In order to avoid the situation where two or more hidden nodes again select overlapping slots, each node in question selects for its message transmission only that slot which is not yet occupied by other nodes and preferably from slots that have been free for a long time, that is to say have not been occupied for a long time, and/or the slot which was last occupied by the node, for which no current occupation exists in the table, that is to say there is no "yes" entry for this node. For example, if slot 4 was used by node 342 in the penultimate frame, which node is currently occupying slot 2, whereas slots 1 and 3 have been free for a long time, the node may select one of these latter slots.

LIST OF REFERENCES:

A, B, C	nodes
1, 2	nodes
r	radio range of a node